

REMARKS / ARGUMENTS

Claims 43-52 remain pending in this application. No claims have been canceled. New claims 47-52 have been added.

Priority

Applicants appreciate the Examiner's acknowledgment of the claim for priority and safe receipt of the priority document.

Information Disclosure Statement

Applicants respectfully request the Examiner initial reference AS, sign and return a copy of the PTO-1449 Form filed on February 28, 2002. As requested by the Examiner in the Office Action, a copy of the Oracle Corp. document is enclosed. A clean copy of the PTO-1449 Form is also enclosed for the Examiner's convenience.

35 U.S.C. §103

Claims 43-46 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Lowenthal et al (U.S. Patent No. 6,035,306) in view of Ledain et al (U.S. Patent No. 6,021,408). These rejections are traversed as follows.

Claims 43-46 are clearly patentable over the cited art. According to the invention recited in these claims, information used for properly relocating storage

areas is acquired from database management system (DBMS) objects such as tables, indexes and logs. Thus, it is possible to relocate storage areas in a precise manner on the basis of these DBMS objects. Therefore, good performance can be obtained reflecting the information regarding tables, indexes and logs on the storage locations of data.

On the other hand, Lowenthal et al disclose a DBMS that acquires information for relocating storage areas from the Oracle TableSpace, as shown in Fig. 5. With respect to the relationship between the Oracle file and the DBMS objects, Lowenthal et al merely state that in the Oracle system, database objects are all stored as Oracle files, regardless of the type of object (see column 6, lines 37-38). Therefore, Lowenthal et al merely disclose that database objects are stored in Oracle files and do not consider or disclose how database objects are stored in the Oracle files. In actuality, in the Oracle DBMS, a plurality of DBMS objects can be allocated to one TableSpace. Therefore, Lowenthal et al clearly do not disclose or suggest determining storage areas on a DBMS object basis.

As a result, Lowenthal et al cannot relocate storage areas on a DBMS object basis. In Fig. 19 of Lowenthal et al, the storage position movement is performed on a Plex basis, where Plex is a unit of storage area management in a storage layer. Lowenthal et al are silent about any mapping information between the Plex and the DBMS objects. A plurality of DBMS objects (data) could be stored in one Plex. Lowenthal et al fail to disclose or suggest any commands that would be used to migrate only the data of a particular DBMS object. In other words, Lowenthal et al

are silent about any commands for moving storage areas of a particular DBMS object.

On the other hand, according to the present invention, data storage position information 555 is acquired from the DBMS (see Figs. 5 and 10). This data storage position information 555 holds mapping information including sets of a data file path name 562 and a file block number 563 on a data structure name 561 basis. The data file path name 562 is an identifier of a file in which the corresponding data is stored and the file block number 563 indicates the storage location of the file identified (see page 45, lines 7-15). It is noted herein that the data structure implies tables, indexes and logs. Therefore, according to the present invention, storage locations (including locations in files) can be obtained for each of tables, indexes or logs on a table, index or log basis.

By utilizing mapping information, data can be migrated between storage areas. More specifically, the storage location of a table/index log is obtained at a resolution of one block in the physical storage device by using data relocation work information 670 or 670b (see Figs. 15 and 29). Thus, a plan of data migration is generated. Thereafter, a data relocation command is issued (see step 2009 in Fig. 13). At this time, the data migration command specifies and area in either a virtual volume provided by virtual volume switch 72 or a volume provided by storage apparatus 10 (see page 56, line 20 to page 57, line 5).

The deficiencies in Lowenthal et al are not overcome by resort to Ledain et al. The Examiner merely relies upon Ledain et al for teaching writing migration of data to

archival storage. Ledain et al do not disclose the features mentioned above that distinguish the present invention from Lowenthal et al.

New claims 47-52 correspond to previously pending claims 14, 18, 19, 34, 38 and 39 contained in the amendment filed September 22, 2004. It is submitted that these claims also patentably define the present invention over the cited art.

According to the present invention, the data relocation method is based on I/O loads as well as based on other aspects (see Figs. 17 and 19-23). As such, data relocation can be optimally carried out before execution of certain processing. Thus, performance problems can be removed in advance.

On the other hand, Lowenthal et al disclose finding a “hot spot” based on load monitoring data (see column 2, lines 46-58) and a data relocation plan is then created (see Fig. 19). However, data relocation optimization cannot be attained without executing processing to collect the load monitoring data. Therefore, if there is a performance problem, data relocation optimization is performed after placing a load on the system by executing processing. On the other hand, as in the present invention, if only a schema of a data is defined and data regions are allocated, data relocation optimization can be attained based on rules of certain aspects.

Although the Examiner has previously alleged that Lowenthal et al’s system can detect potential “hot spots” before they cause performance problems (citing column 3, lines 1-8), it should be clear based on column 2, lines 46-58, that it is assumed that load monitoring data has to be collected to detect potential “hot spots” before they cause performance problems. In other words, this portion suggests that

possible "hot spots" can be detected in the case when performance is gradually deteriorated although the load monitoring data has been collected.

Furthermore, according to Lowenthal et al's system, a problem may not be detected depending upon the sampling period of the load monitoring data. It is necessary to detect a high-activity state on the load monitoring data to detect a "hot spot". Therefore, it is possible that "hot spots" fail to be detected due to the data sampling period. As an example, if I/O loads are concentrated for 5 seconds in a one minute sampling interval, the operation's rate would be calculated to be at a ratio of 5/60 (approximately 8.3%). In this case, it is likely that this 5 second "hot spot" is undetected.

On the other hand, according to the present invention, data relocation optimization is created based on rules of certain aspects without using load monitoring data. Therefore, problems associated with the data sampling period do not occur. As such, it is submitted that the pending claims patentably define the present invention over the cited art.

Request for Interview

Applicants request that the Examiner conduct an interview with the undersigned in order to speed the prosecution of this application. In this regard, the Examiner is hereby invited to contact the undersigned by telephone in order to arrange a suitable time for such interview.

Appl. No. 10/084,540
Amendment dated December 8, 2005
Reply to Office Action of July 12, 2005

ASA-1072

Conclusion

In view of the foregoing, Applicant respectfully requests that a timely Notice of Allowance be issued in this case.

Respectfully submitted,

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.

By 
Shrinath Malur
Reg. No. 34,663
(703) 684-1120